arXiv:1703.10297v1 [math.NA] 30 Mar 2017

A Variable Step Size Implicit-Explicit Scheme for the Solution of the Poisson-Nernst-Planck Equations

David Yan[1], Mary C. Pugh[2], and Francis P. Dawson[1]

[1]Department of Electrical and Computer Engineering, University of Toronto

[2]Department of Mathematics, University of Toronto

March 31, 2017

Abstract

The Poisson-Nernst-Planck equations with generalized Frumkin-Butler-Volmer boundary conditions (PNP-FBV) describe ion transport with Faradaic reactions and have applications in a wide variety of fields. In this article, we develop a variable step size implicit-explicit time stepping scheme for the solution of the PNP-FBV equations. We test our numerical scheme on a simplified "toy" version of the PNP-FBV equations, paying special care to the treatment of the coupled nonlinear terms in the boundary condition. We evaluate various ways of incorporating the boundary condition into the scheme and a method based on ghost points is chosen for its favorable numerical properties compared to the alternatives. In fact, we observe that when the underlying dynamics is one that would have the solutions converge to a steady state solution, the numerical simulation does not result in the time-step sizes growing larger along with the expected convergence. We observe that the time-step sizes threshold at a particular step size. By performing an A-stability analysis we demonstrate that this thresholding does not appear to be due to a stability constraint. Using the developed numerical method, we are able to run simulations with a large range of parameters, including any value of the singular perturbation parameter ϵ.

Keywords. Poisson-Nernst-Planck Equations, Implicit-Explicit Methods, Adaptive Time Stepping

1 Introduction

The Poisson-Nernst-Planck (PNP) equations describe the transport of charged species subject to diffusion and electromigration. They have wide applicability in electrochemistry, and have been used to model a number of different systems, including porous media [1, 2, 3, 4], microelectrodes [5, 6], ion-exchange

membranes [7, 8], electrokinetic phenomena [9, 10, 11], ionic liquids [12, 13], electrochemical thin films [14, 15, 16], fuel cells [17], supercapacitors [18], and many more.

The one-dimensional, nondimensionalized PNP equations for a media with 2 mobile species is

$$\frac{\partial c_\pm}{\partial t} = -\frac{\partial}{\partial x}\left[-\frac{\partial c_\pm}{\partial x} - z_\pm\, c_\pm\, \frac{\partial \phi}{\partial x}\right], \qquad t > 0,\ x \in (0,1), \tag{1}$$

$$-\epsilon^2 \frac{\partial^2 \phi}{\partial x^2} = \frac{1}{2}\left(z_+\, c_+ + z_-\, c_-\right), \qquad x \in (0,1), \tag{2}$$

where $c_\pm$ and $z_\pm$ are the concentration and charge number of the positive/negative ion, ϕ is the potential and ϵ is the ratio of the Debye screening length to the interelectrode width L. The PNP equations are combined with generalized Frumkin-Butler-Volmer (FBV) boundary conditions[1] on each species which describe charge transfer reactions at electrodes. For the simulations performed later in this article, we choose $z_+ = -z_- = 1$ with two no-flux boundary conditions and two FBV boundary conditions for equation (1) :

$$-\left(-\frac{\partial c_-}{\partial x} - c_-\,\frac{\partial \phi}{\partial x}\right)\Bigg|_{x=0} = \left(-\frac{\partial c_-}{\partial x} - c_-\,\frac{\partial \phi}{\partial x}\right)\Bigg|_{x=1} = 0 \tag{3}$$

$$-\left(-\frac{\partial c_+}{\partial x} - c_+\,\frac{\partial \phi}{\partial x}\right)\Bigg|_{x=0} = 4k_{c,a}\, c_+(0,t)\, e^{-0.5\ \Delta\phi_{\text{left}}} - 4\, j_{r,a}\, e^{0.5\ \Delta\phi_{\text{left}}}, \tag{4}$$

$$\left(-\frac{\partial c_+}{\partial x} - c_+\,\frac{\partial \phi}{\partial x}\right)\Bigg|_{x=1} = 4k_{c,c}\, c_+(1,t)\, e^{-0.5\ \Delta\phi_{\text{right}}} - 4\, j_{r,c}\, e^{0.5\ \Delta\phi_{\text{right}}}, \tag{5}$$

where $k_{c,a}$, $k_{c,c}$, $j_{r,a}$, and $j_{r,c}$ are reaction rate parameters; the second part of the subscripts a and c refer to the anode and cathode, respectively . Equations (4)–(5) model the electrodeposition reaction

$$\mathrm{C}^+ + \mathrm{e}^- \rightleftharpoons \mathrm{M} \tag{6}$$

where M represents the electrode material. There is a compact layer of charge, called the Stern layer, that occurs in the electrolyte next to an electrode surface [19, 20]. In equations (4)–(5), $\Delta\phi_{\text{left}}$ and $\Delta\phi_{\text{right}}$ refer to the potential differences across the Stern layers that occur at the anode and cathode respectively. Specifically,

$$\Delta\phi_{\text{left}} = \phi_{\text{anode}} - \phi(0,t) = -\phi(0,t), \qquad \Delta\phi_{\text{right}} = \phi_{\text{cathode}} - \phi(1,t) = v(t) - \phi(1,t) \tag{7}$$

where the potential at the anode has been set to zero and $v(t)$ denotes the potential at the cathode. In

[1] written here, for example, for a reaction involving one ionic species and one electron

addition, the Poisson equation (2) uses a mixed (or Robin) boundary condition [14, 15, 16],

$$-\epsilon\,\delta\,\left.\frac{\partial\phi}{\partial x}\right|_{x=0} = \Delta\phi_{\text{left}} = -\phi(0,t), \tag{8}$$

$$+\epsilon\,\delta\,\left.\frac{\partial\phi}{\partial x}\right|_{x=1} = \Delta\phi_{\text{right}} = v(t) - \phi(1,t), \tag{9}$$

where δ is the ratio of the compact layer thickness to L. Finally, there is an ODE which ensures conservation of electrical current at the electrode [21, 22],

$$-\frac{\epsilon^2}{2}\frac{d}{dt}\phi_x(1,t) = j_{\text{ext}}(t) - \left[k_{c,c}\,c_+\,(1,t)\;e^{-0.5\;\Delta\phi_{\text{right}}} - j_{r,c}\,e^{0.5\;\Delta\phi_{\text{right}}}\right], \tag{10}$$

where $j_{\text{ext}}(t)$ is the current through the device. We refer to the PNP equations with the generalized Frumkin-Butler-Volmer as the PNP-FBV equations.

The device is operated in two regimes — either the current or the voltage at the cathode is externally controlled. If the voltage at the cathode, $v(t)$, is externally controlled then the the PNP-FBV equations (1)–(2) with boundary conditions (3)–(5) and (7)–(9) are numerically solved, determining $c_\pm$ and ϕ. The current is found *a postiori* using equation (10). If the current, $j_{\text{ext}}(t)$, is externally controlled then equation (10) is numerically solved along with the PNP-FBV equations determining $c_\pm$, ϕ, and v simultaneously.

Though many workers in the field have approximated solutions to the PNP-FBV equations using asymptotics [23, 24], a numerical method is needed to obtain a full solution. A key mathematical aspect of the PNP-FBV equations is that the parameter ϵ acts as a singular perturbation to the system and results in the formation of boundary layers [25]. This makes numerical simulation of the PNP-FBV equations especially challenging.

In general, there are several properties of a numerical method that need to be established so as to allow the solver for the Poisson-Nernst-Planck equations to be stable, provide accurate resultsf and converge quickly. The first is a nonuniform mesh in space, which is useful because the potential and concentration distributions near the electrode can change significantly over short distances. The mesh near the boundaries should be finer than the mesh in the middle of the domain, where the concentration and potential are nearly constant. Secondly, many electrochemical systems of interest are subject to sudden changes of forcing, with long periods of relaxation with constant or no forcing in between. This, along with the transient dynamics driven by the initial data, gives the problem more than one timescale.

One of the first contributions to the field of time stepping methods for the Poisson-Nernst-Planck equations, was by Cohen and Cooley [26], who used an explicit time stepping method with predictor-corrector time-step refinement to solve the electroneutral equations with constant current. The next contribution was by Sandifer and Buck [27], who solved a system of equations with time-independent

Nernst-Planck equations with constant current and an implicit, iterative method to step the displacement field. Brumleve and Buck [28] solved the full PNP equations with Chang-Jaffe boundary conditions using Backward Euler time stepping. In Brumleve and Buck's method, time steps were allowed to be variable, but were not adjusted during each step; the method given in the original publication still sees use in more modern work [29]. Murphy et. al. [30] solved the full PNP-FBV equations with adaptive steps by treating the discretized parabolic-elliptic system as a differential-algebraic system of ODE's and algebraic equations, and used a variable-order Gear's method [31]. It is also worth mentioning here the work of Scharfetter and Gummel [32], who gave a numerical method to solve the drift-diffusion equations (an analogue of the PNP equations in semiconductor physics) using Crank-Nicolson time stepping.

One example of a recent work where the time-dependent PNP equations are solved is Soestbergen, Biesheuvel and Bazant [22] (with FBV) who used the commercial finite element software COMSOL, which uses BDF and the generalized-α method [33]. Another is Britz and Strutwolf [34], who simulate a liquid junction using BDF with constant time steps. Britz [35] also outlines various explicit and implicit multistep methods in his book on computational solutions to the PNP equations. Our present work differs from previous attempts (notably Murphy et. al.) by using a splitting method for the parabolic-elliptic system (1)–(2) rather than treating them as a system of DAE's, and differs from other modern computational methods [36] by controlling time steps adaptively.

In this article, we develop and test an adaptive time-stepper for the PNP-FBV equations with error control. This will allow the time-stepper to automatically detect changes in time scales and vary the step size accordingly. The adaptive time-stepper is based on a second-order variable step-size, semi-implicit, backward differentiation formula (VSSBDF2) [37]. The second-order constant step-size semi-implicit, backward differentiation formula (BDF2) is used as a time-stepping scheme in a variety of fields in science, engineering and computational mathematics. Axelsson et. al. [38], for example, used BDF2 with constant steps to solve the time-dependent Navier-Stokes equations. Lecoanet et. al. [39] used a similar method to model combustion equations in stars, and Linde, Persson and Sydow [40] used it to solve the Black-Scholes equation in computational finance.

We also present our findings on using an adaptive multistep time stepping scheme to solve the PNP-FBV equations. First, we applied the scheme to a toy model of the PNP-FBV equations in order to study how error and step size control can be incorporated into the VSSBDF2 scheme for the PNP-FBV equations. We find two issues of interest. The first is the discretization of the nonlinear coupled boundary conditions: we test two ways of implementing these boundary conditions and find that incorporating the flux by applying the PDE at boundary points is numerically more favourable. The second is we observe that when the underlying dynamics is one that would have the solutions converge to a steady state solution, the numerical simulation does not result in the time-step sizes growing larger and larger. In fact, we observe that the time-step sizes threshold at a particular step size. By performing an A-stability

analysis we demonstrate that this thresholding does not appear to be due to a linear stability constraint. We also observe this "thresholding" behavior, in which time steps chosen by the adaptive algorithm are unable to rise above a threshold value, for the PNP-FBV equations. We study this thresholding behaviour with numerical experiments, and find that it is unrelated to the A-stability of the VSSBDF2 scheme. It must be due to a nonlinear instability of some sort. We find that with constant time steps, if the time steps are small enough, the solution converges to steady-state, but if the time steps are too large, the solutions blow up. This transition between converging and blowing up with constant time steps is related to the thresholding behavior we find in the adaptive time stepper.

Eckert et. al. [41] used an adaptive, variable step-size BDF2 algorithm, which is similar to the one we present, to model electroplasticity, but do not report any kind of instability or time step constraint. Rosam, Jimack and Mullis [42] also use an adaptive BDF2 algorithm to study a problem in binary alloy solidification. In their Figure 4, they appear to show time steps relaxing to a constant value (i.e. thresholding), but the reason is not given: they report that it is related to the tolerance set in the adaptive time stepper.

The article is structured as follows. Sections 2.1 and 2.2 present the spatial and temporal discretization, respectively, and in Sections 2.3–2.4 we describe our adaptive stepping and error control algorithm. Section 2.5 presents the splitting method we use to compute the parabolic and elliptic PDEs separately. In Section 3, we present a toy model for the PNP-FBV equations and discuss two methods for discretizing the boundary conditions. In Section 3.1 we present the A-stability analysis for VSSBDF2 time-stepping and discuss the observed thresholding for the toy model. In Section 3.2 we discuss the effect of changing the boundary conditions. In Section 4, we apply the scheme to the PNP-FBV equations and study its performance including the effect of the singular perturbation parameter ϵ.

The numerical simulations in this work were written in MATLAB and run on a Windows PC with 4 cores and 16 GB of RAM.

2 The Numerical Method

2.1 Spatial Discretization

The geometry is divided into a non-uniform mesh, with $x \in [0,1]$ and parameterized via

$$x(s) : [0,1] \to [0,1], \tag{11}$$

so that $x_i := x(s_i)$, $0 = x_1 < x_2 < ... < x_N = 1$ and $dx(i) := x(s_{i+1}) - x(s_i)$ where $s_i = (i-1)\,ds$, $i = 1, 2, ..., N$ and $ds = 1/(N-1)$. The function $x(s)$ may, for example, be piecewise linear with smaller slopes near the endpoints $x = 0, 1$ and a large slope around $x = 1/2$. This would result in a piecewise

uniform mesh that is finer near the endpoints. Alternatively, one might use a logistic function, for example, to create a mesh with smoothly varying nonuniformity as in

$$x(s_i) = \begin{cases} 0 & i = 1 \\ \frac{1}{1+e^{-\gamma(s_i - \frac{1}{2})}} & i = 2, ..., N-1 \\ 1 & i = N. \end{cases} \tag{12}$$

The larger the value of γ, the finer the mesh near the endpoints. In practice, we used a piecewise constant mesh with three or five regions of uniform mesh.

For the discussion of spatial discretization of the PDE (1) and boundary conditions (4–5), we refer to a generic continuity equation $u_t = -\left(J(u, u_x)\right)_x$. The function $u(\cdot, t)$ on the interval is approximated by a vector $\mathbf{u}(t) \in \mathbb{R}^N$ with $u_i(t) \approx u(x_i, t)$. At internal nodes, the flux is approximated using a second-order center-differencing scheme

$$\frac{du_i}{dt} \approx (u_t)\Big|_{x_i} = (J(u, u_x))\Big|_{x_i} \approx \frac{J(u, u_x)\Big|_{x_{i+1/2}} - J(u, u_x)\Big|_{x_{i-1/2}}}{x_{i+1/2} - x_{i-1/2}}, \tag{13}$$

where $x_{i+1/2}$ is the midpoint of $[x_i, x_{i+1}]$ and the approximations

$$u_{i\pm1/2} \approx \frac{u_{i\pm1} + u_i}{2}, \qquad u_x\Big|_{x_{i-1/2}} \approx \frac{u_i - u_{i-1}}{x_i - x_{i-1}}, \qquad \text{and} \qquad u_x\Big|_{x_{i+1/2}} \approx \frac{u_{i+1} - u_i}{x_{i+1} - x_i} \tag{14}$$

are used. At the boundary nodes, we use a three-node, second-order approximation for u_x. For example, the approximation at the left hand boundary is

$$u_x\Big|_{x=0} \approx -\frac{2dx_1 + dx_2}{dx_1\,(dx_1 + dx_2)}u_1 + \frac{dx_1 + dx_2}{dx_1 dx_2}u_2 - \frac{dx_1}{dx_2\,(dx_1 + dx_2)}u_3. \tag{15}$$

For the spatial discretization of (2) with boundary conditions (8)–(9), we refer to a generic Poisson equation, $u_{xx} = f(u)$. In the interior, we again use a three point center-differencing scheme

$$\frac{2u_{i+1}}{dx_i\,(dx_i + dx_{i-1})} - \frac{2u_i}{dx_i\,dx_{i-1}} + \frac{2u_{i-1}}{dx_{i-1}\,(dx_i + dx_{i-1})} \approx u_{xx}\Big|_{x_i} = f(u)\Big|_{x=x_i}. \tag{16}$$

On the boundaries $x = 0, 1$, we use a left or right-handed three point stencil to approximate any first derivatives that might appear in the boundary conditions.

Bazant and coworkers (i.e. in [15] and [16]), used a Chebyshev pseudospectral spatial discretization in their work, where we use a finite difference discretization. We wrote and tested a Chebyshev spectral version of the code using the `chebfun` package [43, 44] and found that, while the spectral and finite difference codes gave nearly identical results, the finite difference code ran orders of magnitude faster

than the spectral code when using the same time stepping scheme.

2.2 Time Stepping

Multistep schemes such as backwards differentiation fomulae, Adams-Bashforth, and Crank-Nicholson methods have long been applied in computational fluid mechanics to time step advection-diffusion equations (see Chapter 4.4 of [45]) where both the diffusion and advection terms are linear. Semi-implicit, or implicit-explicit schemes, are useful when PDEs contain both linear stiff terms and nonlinear terms which are difficult to handle using implicit methods. For our model, we use a second-order variable step-size implicit-explicit backwards differencing formula, VSSBDF2, introduced by Wang and Ruuth [37, 46]. To present the VSSBDF2 method, we consider $\frac{du}{dt} = f(u) + g(u)$ where $f(u)$ is a nonlinear term (that is best handled explictly) and $g(u)$ is a stiff linear term that is best handled implicitly to avoid numerical instability.

We use a superscript notation to denote time levels: u^n approximates $u(t^n)$. Given the values u^{n-1} at time $t^{n-1} = t^n - dt_{\text{old}}$ and u^n at time t^n, the future value u^{n+1} at time $t^{n+1} = t^n + dt_{\text{now}}$ is determined by equation 2.8 in [37]:

$$\frac{1}{dt_{\text{now}}}\left(\frac{1+2\omega}{1+\omega}u^{n+1} - (1+\omega)u^n + \frac{\omega^2}{1+\omega}u^{n-1}\right) = (1+\omega)f(u^n) - \omega f(u^{n-1}) + g(u^{n+1}), \qquad (17)$$

where $\omega = dt_{\text{now}}/dt_{\text{old}}$. This scheme was shown to have favorable stability properties compared to the other IMEX and VSIMEX schemes based on numerical experiments on Burger's equation [37].

2.3 Error Approximation and Extrapolation

If one is using a non-adaptive time stepper then the time steps t^n, $n = 0, 1, \ldots$ are chosen before computing u^n, $n = 0, 1, \ldots$, An adaptive time stepper chooses the time step size dt_{now} based on already-computed values. For example, using the values u^{n-1}, $u^{n-\frac{1}{2}}$ and u^n at times $t^n - dt_{\text{old}}$, $t_n - dt_{\text{old}}/2$ and t_n, respectively, we choose a time step size dt_{now} so that the local truncation error (LTE) is "small enough" but not "too small". This is done as follows. First, we choose a candidate time step: $dt_{\text{now}} = dt_{\text{old}}$, for example. We then take one "coarse" step from $t^n - dt_{\text{old}}$ and t^n to $t^n + dt_{\text{now}}$, using u^{n-1} and u^n to create u_c^{n+1}. Next, we take one "fine" step from $t^n - dt_{\text{old}}/2$ and t^n to $t^n + dt_{\text{now}}/2$, using $u_f^{n-1/2}$ and u^n to create $u_f^{n+1/2}$ and take a second fine step from t^n and $t^n + dt_{\text{now}}/2$, using u^n and $u_f^{n+1/2}$ to create u_f^{n+1}. We use u_c^{n+1} and u_f^{n+1} to estimate the local truncation error; if the error is acceptable, then we advance in time. If the error is unacceptable, then we choose a new dt_{now} and try again creating a new u_c^{n+1} and u_f^{n+1} pair. Values of u at half steps are stored at each step in order to allow the next fine step. This scheme is shown in Figure 1. Note that in the figure, $dt_{\text{now}} > dt_{\text{old}}$; this would happen in the first attempt if the local truncation error was too small.

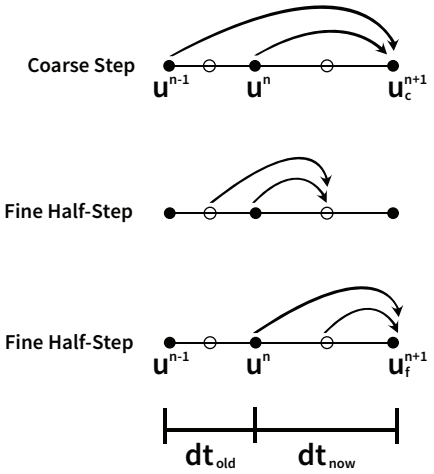

Figure 1: Coarse/fine time stepping scheme. First a coarse time step uses u^{n-1} and u^n to create u_c^{n+1}. Then a fine time step uses $u_f^{n-\frac{1}{2}}$ and u^n to create $u_f^{n+\frac{1}{2}}$ and another fine time step uses u^n and $u_f^{n+\frac{1}{2}}$ to create u_f^{n+1} In the diagram, open circles indicate values from fine half-steps.

Given the two approximations u_c^{n+1} and u_f^{n+1} of the solution at time $t^n + dt_{\text{now}}$, one can use them to approximate the local truncation error (see Appendix A):

$$\epsilon_c = u_c^{n+1} - u(t^{n+1}) \approx \frac{8\,(dt_{\text{old}} + dt_{\text{now}})}{7dt_{\text{old}} + 5dt_{\text{now}}} \left(u_c^{n+1} - u_f^{n+1}\right). \tag{18}$$

We use this approximation of ϵ_c when testing whether to use dt_{now} or to coarsen or refine it. If dt_{now} has been accepted, we then use u_f^{n+1} and u_c^{n+1} to construct an approximation u^{n+1} which has smaller truncation error (see Appendix A). Specifically, u^{n+1} is a linear combination of the form

$$u^{n+1} = \alpha u_c^{n+1} + \beta u_f^{n+1} \quad \text{with coefficients} \quad \alpha = -\frac{dt_{\text{old}} + 3\,dt_{\text{now}}}{7\,dt_{\text{old}} + 5\,dt_{\text{now}}}, \quad \beta = 8\,\frac{dt_{\text{old}} + dt_{\text{now}}}{7\,dt_{\text{old}} + 5\,dt_{\text{now}}}. \tag{19}$$

The local truncation error for u^{n+1} is one order higher than the local truncation errors for u_c^{n+1} and u_f^{n+1}. Note that if $dt_{\text{now}} = dt_{\text{old}}$ then equation (19) reduces to the classic Richardson extrapolation formula.

2.4 Adaptive Time Stepping

We now discuss how to use the approximate local truncation error (18) in the adaptive time stepping method. There are two parameters in the adaptive scheme: *tol* and *range*. The adaptive scheme seeks a time step dt_{now} such that the local truncation error $\epsilon_c^{n+1} \in (tol - range, tol + range)$. In practice, we have used a relationship such as $range = tol/3$ and $tol = 10^{-6}$.

1. Set the loop counter, i, to 1 and set $dt^1 = dt_{\text{old}}$, where dt^i is the i$^{\text{th}}$ guess at a value for dt_{now}.

2. (a) Take coarse and fine steps to compute u_c^{n+1} and u_f^{n+1} and use (18) to approximate ϵ_c^i. Use the multistep scheme for $n > 1$ and equation (17) with $\omega = 0$ (single-step) for the first time step.

   (b) Compute ϵ_c^i and check if $|\epsilon_c^i - tol| < range$. If yes, go to step 5.

3. Check the loop counter. If $i_{\max}$ attempts have been made trying to find an acceptable value for dt_{now} and $|\epsilon_c^i| > tol + range$, then set dt_{now} to $dt_{\min}$. Alternatively, if at any time $dt_{\text{now}} > dt_{\max}$, then set dt_{now} to $dt_{\max}$. Compute the resulting u_c^{n+1} and u_f^{n+1}, and go to step 5.

4. Increment the loop counter by 1 and make a guess at dt_{now} based on equation (20) with $p = 3$ (the order of the local truncation error) and return to step 2

$$dt^{i+1} = dt^i \min\left(\max\left(\left(\frac{tol}{\epsilon_c}\right)^{1/p}, \eta_{\min}\right), \eta_{\max}\right) \tag{20}$$

5. Advance in time, setting $dt_{\text{now}} = dt^i$, $t^{n+1} = t^n + dt_{\text{now}}$ and use (19) to define u^{n+1}.

$i_{\max}$ is the maximum number of allowed iterations per time step, and $dt_{\min}$ and $dt_{\max}$ are the minimum and maximum allowed time steps, respectively. $dt_{\min}$ is set so that the simulation completes in a reasonable period of time; we used values in the range 10^{-10} to 10^{-8} for runs up to around time 50.

Equation (20) is a carry-over from refinement strategies for single step schemes, where p is the order of the LTE. It is based on the idea that the LTE satisfies $\epsilon^i \approx C\left(dt^i\right)^p$, so setting $dt^{i+1} = \left(\frac{tol}{\epsilon^i}\right)^{1/p} dt^i$ will bring ϵ^{i+1} closer to tol. The role of $\eta_{\min}$ and $\eta_{\max}$ in (20) is to prevent dt^{i+1} from changing "too much" at once ($\eta_{\min} \leq dt^{i+1}/dt^i \leq \eta_{\max}$). For $\eta_{\max}$ and $\eta_{\min}$, we used values such as 1.2 and 0.8, respectively, as safeguards to guarantee that the adaptive stepper is able to converge. This type of error control strategy is discussed in Chapter II.4 of Hairer, Norsett and Wanner [47]. Since we're using a two-step time stepping scheme, for the first time step, we use a first order IMEX scheme, i.e. (17) with $\omega = 0$, which has an LTE $\sim O(dt^2)$ along with the error control that uses $p = 2$ in (20) and the extrapolation formula $u^{n+1} = 2u_f - u_c$.

Above, we described the time-stepping for a single ODE according to the steps presented above. Given the system of ODEs that would arise from spatially discretizing a parabolic PDE, we use the same approach but define the local truncation error using the norm of $\mathbf{u}$. In our simulations, we chose the l^2 norm.

2.5 Splitting Method for the Parabolic-Elliptic System

The model presented in the Introduction has 2 parabolic PDEs (equation (1)) for the two species $c_\pm$, one ODE (equation (10)) for the electric field at the cathode ϕ_x (if there are current boundary conditions), and one elliptic PDE (equation (2)) for the potential ϕ . If the model is extended to include adsorption effects at the electrodes, there would be an additional ODE describing the dynamics of the fraction of surface coverage for each electrode [48]. If, in addition, the model included temperature and heat transport, there would be another PDE for the temperature [48]. In this work however, we limit ourselves to just the PNP-FBV system of equations.

We now describe how to time-step this system in a way that the global truncation error is $O(dt^2)$ when we take constant time steps ($dt_{\text{old}} = dt_{\text{now}} = dt$). Figure 2 shows the scheme for the first and

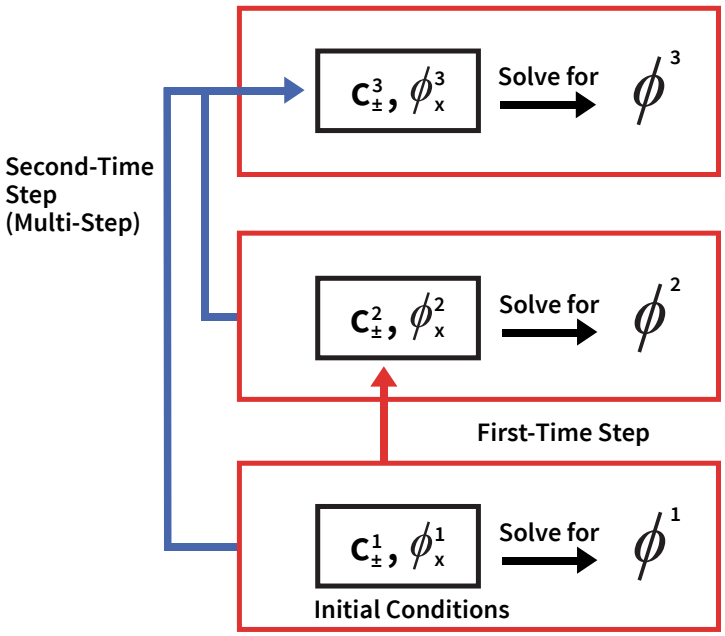

Figure 2: A diagram of the first three time levels. $\mathbf{c}^1_\pm$ and ϕ^1_x are the initial conditions, which are used by the elliptic solver to generate for $\boldsymbol{\phi}^1$. The four vectors at the first time are then all used to make a step using a first-order IMEX scheme, creating $\mathbf{c}^2_\pm$ and ϕ^2_x. These three vectors are then used to generate $\boldsymbol{\phi}^2$, and all eight vectors from the first and second time levels are used in a two-step scheme to create $\mathbf{c}^3_\pm$ and ϕ^3_x.

second (and subsequent) time steps. The initial value $\phi_x(1,0)$ is denoted ϕ^1_x. The initial data $c_\pm(x,0)$ is discretized resulting in initial vectors $\mathbf{c}^1_+, \mathbf{c}^1_- \in \mathbb{R}^N$. Given $\mathbf{c}^1_\pm$ and ϕ^1_x, we solve the elliptic PDE to determine $\boldsymbol{\phi}^1 \in \mathbb{R}^N$. Using a first-order IMEX scheme on the two parabolic PDEs (1) and the ODE (10), we determine $\mathbf{c}^2_\pm$ and ϕ^2_x. We then solve the elliptic PDE (2) to determine $\boldsymbol{\phi}^2$. The two-step method is then applied to determine $\mathbf{c}^3_\pm$ and ϕ^3_x (and subsequent values of $\mathbf{c}^n_\pm$ and ϕ^n_x). In general, $\boldsymbol{\phi}^{n+1}$ is solved for using $\mathbf{c}^{n+1}_\pm$ and ϕ^{n+1}_x. It is important for numerical accuracy that the correct $\boldsymbol{\phi}$ is associated with the correct time level: not doing so reduces the global truncation error from $O(dt^2)$ to $O(dt)$. Furthermore, because ϕ at time t^n is determined from the other quantities at time t^n the elliptic equation does not factor into the computation of the local truncation error during adaptive time stepping.

3 Numerical Tests on Model Problem

We begin by exploring the numerical properties of the time-stepping scheme for a simple coupled parabolic-elliptic system which has many of the properties of the full system of equations (1)–(9). We consider the

"toy" model

$$u_t = u_{xx} + (uv_x)_x = -(-u_x - u\,v_x)_x, \tag{21}$$

$$v_{xx} = -u \tag{22}$$

on $x \in (0,1)$ with boundary conditions on the flux of u

$$(-u_x - uv_x)\Big|_{x=0} = F(t) := -u(0,t)\, e^{v(0,t)} + e^{-v(0,t)}, \tag{23}$$

$$(-u_x - uv_x)\Big|_{x=1} = G(t) := u(1,t) - 1. \tag{24}$$

We refer to the functions $F(t)$ and $G(t)$ as the "flux constraint functions". Finally, there are Robin and Neumann boundary conditions on v

$$v(0,t) - 0 = v_x(0,t), \tag{25}$$

$$v_x(1,t) = 0. \tag{26}$$

This system shares two important properties with the PNP-FBV equations. First, the flux of u, $-u_x - uv_x$, is nonlinear and coupled to v. Second, the flux constraint function, $F(t)$ in (23), is nonlinear and coupled to v. We refer to this boundary condition as "Butler-Volmer-like", or "BV-like".

In the interior of the domain, $x \in (0,1)$, we discretize in space using the methods outlined in Section 2.1 and time step using the time stepper, error control, and splitting methods outlined in Sections 2.2–2.5. At the interior nodes, $x_2, \ldots x_{N-1}$, the PDE (21) is temporally discretized using the VSSBDF2 scheme (17). The term u_{xx} in (21) is handled implicitly, like the term $g(u)$, and the nonlinear term $(u\,v_x)_x$ is extrapolated forward to time t^{n+1}, like the nonlinear term $f(u)$. This results in $N-2$ equations involving the N unknowns, $u_1^{n+1}, \ldots, u_N^{n+1}$. The boundary conditions provide two additional equations.

Above, we have described a strategy for updating the interior points in time. However, it is unclear how the boundary conditions on the flux of u, equations (23)–(24), should be handled. A natural first approach is to extrapolate both the nonlinear term in the flux and the flux constraint function forward to time t^{n+1}, using the equation

$$-(u_x)_1^{n+1} - (1+\omega)(uv_x)_1^n + \omega(uv_x)_1^{n-1} = (1+\omega)\,F^n - \omega\,F^{n-1} \tag{27}$$

to approximate the boundary condition (23), for example. Here, $F^n = -u_1^n\, e^{v_1^n} + e^{-v_1^n}$; F^{n-1} is defined analogously. Similarly, the boundary condition (24) would be approximated as

$$-(u_x)_N^{n+1} - (1+\omega)(uv_x)_N^n + \omega(uv_x)_N^{n-1} = (1+\omega)\,G^n - \omega\,G^{n-1} \tag{28}$$

where $G^n = u_N^n - 1$ and $G^{n-1} = u_N^{n-1} - 1$. We refer to equations (27) and (28) as the "direct" method of handling the boundary conditions. The N equations are then solved, thus determining $u_1^{n+1}, \ldots u_N^{n+1}$. The elliptic equation $v_{xx} = -u^{n+1}$ and the Robin and Neumann boundary conditions (25)–(26) are then discretized and solved, thus determining $v_1^{n+1}, \ldots v_N^{n+1}$.

As shown in Table 1, with constant time steps, we found the "direct" method to be second-order accurate in time if the boundary conditions on the flux of u are replaced with homogeneous Dirichlet boundary conditions on u: $u(0,t) = u(1,t) = 0$. The method is first-order accurate in time with the BV-like boundary conditions, (27)–(28). Because the BV-like boundary conditions are nonlinear and coupled to the elliptic equation, we also considered boundary conditions of intermediate complexity. Specifically, the linear, coupled boundary conditions $u(0,t) = 0$ and $u(1,t) = v(1,t) - 1$ and the nonlinear, uncoupled boundary conditions $u(0,t) = 0$ and $u(1,t) = \cos(u(1,t)) - 1$. The method was found to be unstable with no-flux boundary conditions on u: $F(t) = G(t) = 0$ in (23)–(24).

dt	Homog. Dirichlet	linear coupled	nonlinear uncoupled	BV-Like
0.001	4.01	3.58	2.05	2.43
0.001/2	4.01	2.89	2.02	2.30
0.001/4	4.01	4.79	2.01	2.18
0.001/8	4.00	11.26	2.01	2.10
0.001/16	4.02	1.14	2.00	2.06

Table 1: Convergence test on toy model The boundary conditions are: homogeneous Dirichlet: $u(0,t) = u(1,t) = 0$, linear & coupled: $u(0,t) = 0$ $u(1,t) = v(1,t) - 1$, nonlinear & uncoupled: $u(0,t) = 0$ $u(1,t) = \cos(u(1,t)) - 1$, and BV-like: (27)–(28). We compute six solutions on a uniform mesh $dx = 1/20$ using constant time steps up to time $t = 1$. The i^{th} solution $\mathbf{u}_i$ is computed using $dt = 0.001/2^i$. The ratios are computed using the discrete solutions $\mathbf{u}_i$ at the final time $t = 1$; the ratio is defined as $\|\mathbf{u}_i - \mathbf{u}_{i+1}\|/\|\mathbf{u}_{i+1} - \mathbf{u}_{i+2}\|$ with the l^2 norm. Ratios approaching 4 indicate second order accuracy in time, while ratios approaching 2 indicate first order accuracy. Note that, while we only show 5 values of dt in the table, additional runs with $dt = 0.001/32$ and $dt = 0.001/64$ were necessary to create the ratios.

We found that with BV-like boundary conditions, the error control algorithm is unable to make the error arbitrarily small. Specifically, the adaptive time-stepper chooses the time steps so that the local truncation error is in the interval $(tol - range, tol + range)$ and we found that we could not take tol to be smaller than .009.

If we take constant, small time steps, the simulation shows a smooth solution converging exponentially fast to a steady state solution: u_t and the higher temporal derivatives are decaying exponentially fast to zero. The local truncation error (41) is proportional to $dt_{\mathrm{now}}^2(dt_{\mathrm{now}} + dt_{\mathrm{old}})$ times a linear combination of u_t, u_{tt}, and u_{ttt}. As a result, if constant time steps are used the local truncation error should decay exponentially fast to zero; this is what we observe in our simulations with BDF2.

Based on this observation that the initial value problem has a solution that converges to a steady state, one would make the following prediction about what would happen when an adaptive time-stepper is used. Because u_t and the higher temporal derivatives are decaying exponentially fast to zero and the local truncation error is constrained to stay within a fixed interval, $[tol - range, tol + range]$ with

$tol - range > 0$, one would expect that the time-step would grow in order to satisfy this constraint. The time-step would increase until it reaches the maximum time-step size $dt_{\max}$ (see Subsection 2.4) at which point it would stay at this maximum time-step size. Once dt reaches the maximum value $dt_{\max}$ the local truncation error is no longer forced to be greater than $tol - range$. At this point in time, the local truncation error would decrease to zero exponentially fast.

In fact, this is precisely what one observes if one computes solutions of the diffusion equation with homogeneous Dirichlet or Neumann boundary conditions using either Backward Euler timestepping or the VSSBDF2 timestepping (see Figure 3.3 in [48]). In contrast, when using VSSBDF2 to compute solutions of the system (21)–(26) using the direct method (27)–(28) for the boundary conditions, we find that initially the solution starts to converge to the steady state and the time-step sizes grow in order to keep the local truncation error within the specified bounds. The time-step sizes do not increase to $dt_{\max}$ though — instead they "threshold" at some smaller value of dt. Further, u_t and the higher derivatives stop decreasing in magnitude. We find that the threshold, dt_∞, does not depend on the mesh size when using a uniform mesh on $(0, 1)$.

The "direct method" for boundary conditions (27)-(28) has the undesirable properties of: the method is not second-order accurate in time for the BV-like boundary conditions, we cannot take the parameter tol to be arbitrarily small, and the time-step dt thresholds. For this reason, we used an alternative approach for the boundary conditions. We refer to it as the "ghost point" method (see Section 1.4 of Thomas [49]), since it assumes the PDE (21) holds at the two boundaries, $x = 0, 1$, and uses the same time stepping scheme at the boundaries as at the internal nodes $x_2, \ldots x_{N-1}$. Applying the PDE at x_1 and x_N requires the flux at neighbouring points; in a true ghost point method this flux would be located at "ghost points" outside the computational domain. Instead, we use the flux *at* x_1 and x_N; there are no points outside the computational domain. The spatially discretized PDE at $x = 0, 1$ becomes the two ODEs

$$u_t\Big|_{x=0} = \frac{(u_x + uv_x)\Big|_{x=dx_1/2} - F(t)}{dx_1/2} \tag{29}$$

$$u_t\Big|_{x=1} = \frac{G(t) - (u_x + uv_x)\Big|_{x=1-dx_{N-1}/2}}{dx_{N-1}/2} \tag{30}$$

As shown in Table 2 , when using BDF2 timestepping we found the ghost point method to be second-order accurate in time for BV-like boundary conditions (23)-(24). This is as expected because the time stepping scheme is second-order accurate and the ghost point method steps the boundary points in the same manner as the interior points.

The ghost point method also performs better with the adaptive time stepper in that the local truncation error can be made arbitrarily small by decreasing dt: we find that any value of tol above round-off is

dt	Ratios
0.001	3.47
0.001/2	3.65
0.001/4	3.79
0.001/8	3.89
0.001/16	3.94

Table 2: Convergence Test on Toy Model with ghost point method and BV-like boundary conditions. See the caption of Table 1 for a full explanation. The mesh width is $dx = 1/20$ and the simulation ran to time $t = 1$. Ratios tending to 4 indicate second-order accuracy in time. Runs with $dt = 0.001/32$ and $dt = 0.001/64$ were necessary to create the ratios.

admissible. However, the method does still have the thresholding behavior with adaptive time stepping and the threshold, dt_∞, does not depend on the mesh size when using a uniform mesh on $(0, 1)$, nor do they depend on the value of tol. The left plot of Figure 3 presents data from a BDF2 simulation. The top plot shows that $\|u_t\|$ is decaying exponentially fast to zero. The middle plot shows that the local truncation error is also decaying exponentially fast to zero. The bottom plot shows the time-step size as a function of time. The right plot presents data from the adaptive time-stepper simulation based on the VSSBDF2 scheme. The top plot shows that $\|u_t\|$ decays exponentially fast, at first, but then plateaus at around 10^{-7}. The middle plot shows that the the local truncation error is in $(tol - range, tol + range)$; this is as expected as long as the time-step size is less than the maximum time-step size $dt_{\max}$. The bottom plot shows that the time-step size has thresholded at a value of around $dt_\infty = .25$; this is less than $dt_{\max} = 1$.

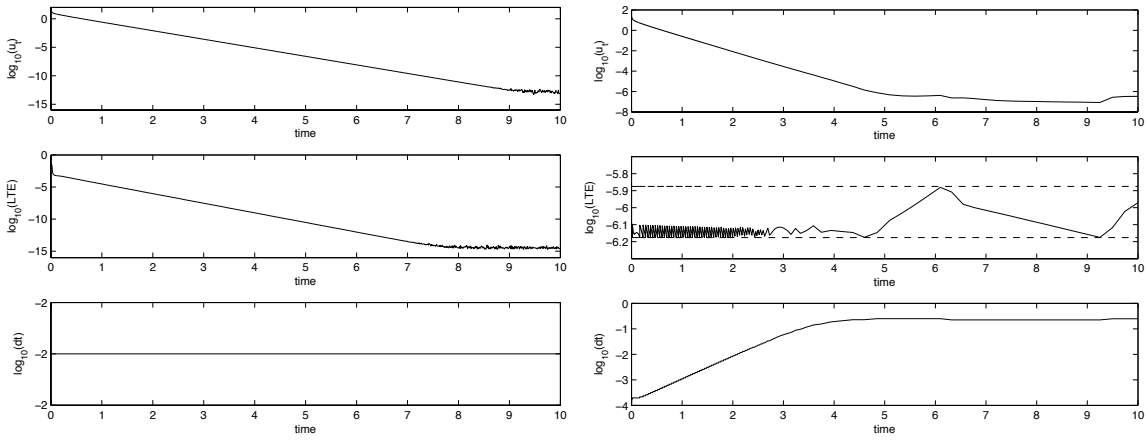

Figure 3: Toy model with BV-like boundary conditions. $N = 40$ and the initial data is $u(\cdot, 0) = 1$. The figure on the left presents data from a simulation with BDF2 time-stepping with $dt = 0.01$. The figure on the right presents data from the adaptive scheme that uses VSSBDF2 time-stepping with $tol = 10^{-6}$ and $range = tol/3$. In both figures, the top plot is a plot of $\log_{10}(\|u_t(\cdot, t)\|)$ versus time. The middle plot is a plot of $\log_{10}(LTE)$ versus time where LTE is the local truncation error. The bottom plot is a plot of $\log_{10}(dt)$ versus time.

It is natural to ask whether the thresholding behaviour still occurs with adaptive time stepping on a one-step time-stepping scheme. We considered a pair of one-step schemes for the toy model PDE (21):

Forward Euler,

$$\frac{u^{n+1}-u^n}{dt} = \left(u_{xx} + (uv_x)_x\right)^n , \tag{31}$$

and the implicit-explicit scheme,

$$\frac{u^{n+1}-u^n}{dt} = (u_{xx})^{n+1} + \left((uv_x)_x\right)^n . \tag{32}$$

For both time-steppers, we tried both the "direct" and ghost point methods for the boundary conditions, (27)-(28) and (29)-(30) respectively. Adaptive stepping is used, with the error at each step calculated via. $\epsilon = |u_{\text{fine}} - u_{\text{coarse}}|$. All four schemes display the same behaviour as seen in the right plots of Figure 3: the time-step thresholds at a value less than $dt_{\max}$ (See Figure 3.9 in [48]). Additionally, the two schemes which used the direct method for the BV-like boundary conditions share the issue of a hard lower bound on the error (*tol* can't be taken small) as was the case for the two-step direct scheme. For this reason, the implicit-explicit method with the ghost point approach to the boundary conditions was used to take the first time step before proceeding with the two-step VSSBDF2 scheme.

3.1 Adaptive Time Stepping and A-stability

A natural question is whether the time-steps are thresholding because there is a linear stability constraint on the time-step size. We consider A-stability; it considers how the scheme performs for the ODE $u_t = \lambda\, u$ where $Re(\lambda) < 0$.

To see how the adaptive time-stepper works for a time-stepping scheme that has stability constraints, we consider Forward Euler time-stepping on the diffusion equation $u_t = u_{xx}$ with homogenous Dirichlet boundary conditions $u(0,t) = u(1,t) = 0$. The A-stability analysis yields $u^{n+1} = \rho\, u^n = (1 + dt\,\lambda)\, u^n$. Forward Euler is A-stable if $|\rho| \leq 1$; that is if the time-step satisfies $dt < -2Re(\lambda)/|\lambda|^2$. For the diffusion equation, one approximates the Laplacian and the boundary conditions using a matrix and then takes λ to be the negative eigenvalue with largest magnitude. Specifically, a three-point finite-difference approximation of the Laplacian on a uniform mesh yields a time-step A-stability constraint of

$$dt < \frac{dx^2}{1+\cos(dx\,\pi)} = \widehat{dt}, \qquad \widehat{dt} = \frac{dx^2}{2} + \mathcal{O}(dx^4). \tag{33}$$

The left plot of Figure 4 presents data from the adaptive time-stepper with Forward Euler time-stepping. The top plot presents the local truncation error as a function of time; it stays in the interval $(tol - range, tol + range)$. The middle plot presents the time-step as a function of time; it thresholds around the value 3.310^{-4}. This limiting value was always less than the A-stability constraint $\widehat{dt}$ (33) and, of course, less than the maximum allowed time step size $dt_{\max}$. The bottom plot presents the root ρ as a function of time; it's roughly at the stability threshold of $|\rho| = 1$ at most times. We repeated the

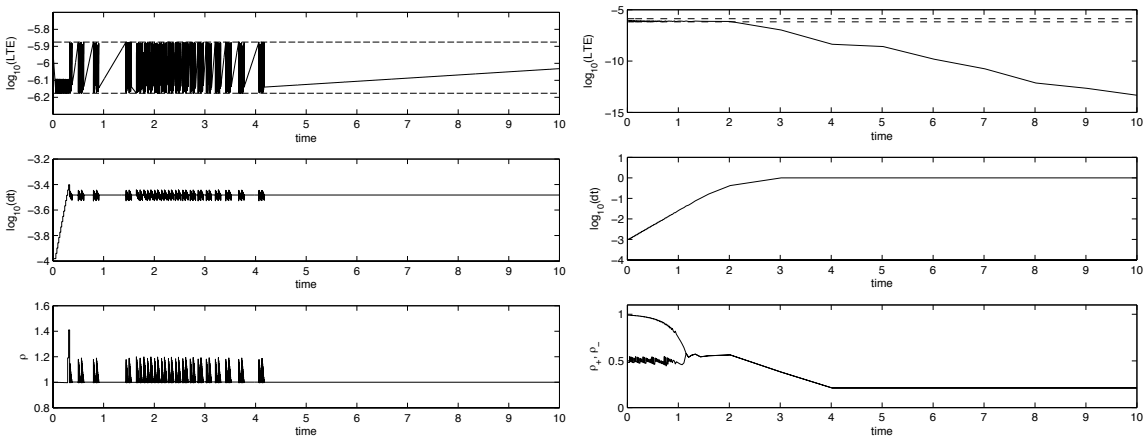

Figure 4: Diffusion equation $u_t = u_{xx}$ with Dirichlet boundary conditions $u(0,t) = u(1,t) = 0$. $N = 40$ and the initial data is $u(x,0) = \sin(\pi\, x)$. The figure on the left presents data from the adaptive scheme using Forward Euler time-stepping. The figure on the right presents data from the adaptive scheme using VSSBDF2 time-stepping. In both cases, $tol = 10^{-6}$, $range = tol/3$, and $dt_{\max} = 1$. In both figures, the top plot is a plot of $\log_{10}(LTE)$ versus time where LTE is the local truncation error. In both figures, the middle plot is a plot of $\log_{10}(dt)$ versus time. In the left figure, the bottom plot is a plot of the largest root ρ from the A-stability analysis for Forward Euler. In the right figure, the bottom plot is a plot of the two largest roots $\rho_\pm$ from the A-stability analysis for VSSBDF2.

simulation for various meshwidths and always found the limiting time-step size, dt_∞, to be less than the A-stability constraint on the time-step size [48].

We now consider the A-stability of the VSSBDF2 scheme given by (17) for the ODE $u_t = \lambda\, u$. In applying the scheme, we have no nonlinear term ($f(u) = 0$) and the linear term ($g(u) = \lambda\, u$) is handled implicitly. The two-step scheme has the general solution $u^n = c_1\rho_+^n + c_2\rho_-^n$ where $\rho_\pm$ are the roots of the quadratic equation

$$\frac{1+2\omega}{1+\omega}\rho^2 - (1+\omega)\,\rho + \frac{\omega^2}{1+\omega} = \lambda\, dt_{\text{now}}\,\rho^2.$$

The roots are given by

$$\rho_\pm = \frac{1}{2\left(\frac{1+2\omega}{1+\omega} - \lambda\, dt_{\text{now}}\right)}\left(1+\omega \pm \sqrt{(1+\omega)^2 - \frac{4\omega^2}{1+\omega}\left(\frac{1+2\omega}{1+\omega} - \lambda\, dt_{\text{now}}\right)}\right). \tag{34}$$

Recalling that ω is the time-step ratio ($\omega = dt_{\text{now}}/dt_{\text{old}}$), one can easily verify that ρ_+ is the essential root and ρ_- is the spurious root: $\rho_+ \to 1$ and $\rho_- \to 0$ as $dt_{\text{now}} \to 0$.

In practice, the adaptive time-stepper has already determined dt_{old} and the differential equation determines the value of λ. If λ is nonzero and real-valued then

$$\rho_\pm \to -\frac{1}{2\,\lambda\, dt_{\text{old}}}\left(1 \pm \sqrt{1+4\,\lambda\, dt_{\text{old}}}\right) \qquad \text{as } dt_{\text{now}} \to \infty. \tag{35}$$

If either of the limiting values has magnitude greater than one, then this will result in an A-stability constraint on dt_{now}; if dt_{now} is too large then at least one of the roots will have a magnitude greater than

one. For example, if $\lambda\, dt_{\text{old}} = -1/8$, then the limiting values are $4 \pm 2\sqrt{2}$; in such a case there would be an A-stability constraint on dt_{now}.

For $\lambda = 0$, A-stability is the same as zero-stability. In this case, there are two roots:

$$\rho_+ = 1 \qquad \text{and} \qquad \rho_- = \frac{\omega^2}{1 + 2\,\omega}. \tag{36}$$

The essential root ρ_+ corresponds to an exact solution of the ODE: $u(t) = c$. As $dt_{\text{now}} \to \infty$, the spurious root ρ_- increases to infinity; thus the spurious root can be a source of numerical instability. Indeed [37], the time-step constraint for zero-stability is $dt_{\text{now}} \leq (1 + \sqrt{2})\, dt_{\text{old}}$.

From our study on Forward Euler, the A-stability constraint on the time-step depended on the ODE or PDE only (via λ). The A-stability constraint on the time-step size stopped the time steps from growing arbitrarily large; there was a limiting time-step size. However, for the VSSBDF2 scheme whether or not a stability constraint comes into play depends on the value of dt_{old}, and so one could imagine that if dt_{old} and dt_{new} could both grow large while the magnitudes of the roots, $|\rho_\pm|$, remains less than 1. For example, using formula (34) one can check that if the time-steps are growing exponentially in time, $t^n \sim \exp(\mu n)$, then the roots $\rho_\pm$ will decay like $1/\sqrt{t^n}$ as $n \to \infty$. We find a similar $1/\sqrt{t^n}$ rate of decay if the time steps are growing algebraically in time ($t^n \sim n^\alpha$ for some fixed $\alpha > 0$).

The right plot of Figure 4 presents results from the adaptive scheme with the VSSBDF2 time-stepper applied to the diffusion equation problem. The top plot shows that that the local truncation error is initially in the interval $(tol - range, tol + range)$ but, after a certain time, it begins to decrease exponentially fast. The middle plot shows that the time-step size increases until it reaches the maximum time-step size, $dt_{\max}$. It is precisely when the time-step is held at $dt_{\max}$ that the local truncation error is no longer constrained to be in the interval and it starts to decay exponentially. The simulation was done with $N = 40$ meshpoints and so the Laplacian with Dirichlet boundary conditions are approximated by a 38×38 matrix. Its 38 eigenvalues result in 38 values for ρ_+ and 38 values for ρ_-. The bottom plot presents the two roots of largest magnitude as a function of time; the magnitudes are always less than 1.

This demonstrates that although, in principle, there could be an A-stability constraint on the time-step size for VSSBDF2, it does not come into play for the diffusion equation. A natural question is whether it might come into play for the toy model with BV-like boundary conditions. The top-left plot of Figure 5 presents the time-step size versus time; it shows the thresholding time-step to be around $dt = 0.25$. If $N = 40$ then the Laplacian matrix for the ghost-point approach is a 40×40 matrix; the resulting 40 eigenvalues yield 40 essential roots, ρ_+, and 40 spurious roots, ρ_-. The bottom-left plot of the figure presents the six largest-magnitude roots. We note that one of the roots has magnitude equal to 1. This arises because when time-stepping the PDEs at the boundary, (29)–(30), the terms u_x are handled implicitly, just like the terms u_{xx} at the internal points. As a result, the Laplacian matrix comes

from discretizing the Laplacian with Neumann boundary conditions; this yields an eigenvalue $\lambda = 0$. The zero-stability analysis (36) gives the essential root $\rho_+ = 1$.

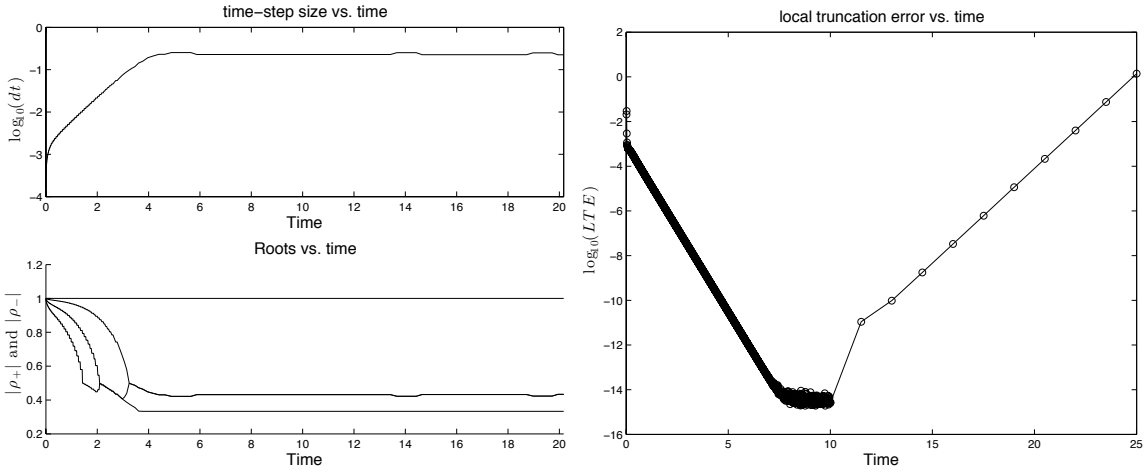

Figure 5: Toy model with BV-like boundary conditions implemented with "ghost point" method. $N = 40$ and $u(x, 0) = 1$. Left plot: the adaptive scheme is used with VSSBDF2 with $tol = 10^{-6}$, $range = tol/3$, and $dt_{\max} = 1$. The top plot is a plot of $\log_{10}(dt)$ versus time. The bottom plot is a plot of the magnitudes of the six largest roots $\rho_\pm$ versus time. (There are two pairs of complex-conjugate roots; they have equal magnitude.) Right plot: An example of piecewise constant time-stepping. The time stepping is done using BDF2 with $dt = 0.01$ for $t \leq 10$ and $dt = 1.5$ for $t > 10$. The plot shows $\log_{10}(LTE)$ versus time where LTE denotes the local truncation error. Note that $dt = 0.01$ is less than the limiting value of dt in the left plot and $dt = 1.5$ is larger than it.

The thresholding of the time-steps causes $\omega \to 1$ as $dt \to dt_\infty$. For this reason, the roots $\rho_\pm$ will converge to those of the A-stability for BDF2. Taking $\omega = 1$ in (34) yields the roots

$$\rho_\pm = \frac{1}{3/2 - \sigma}\left(1 \pm \frac{1}{\sqrt{2}}\sqrt{\sigma + 1/2}\right)$$

where $\sigma = \lambda\, dt$ and we have assumed $\lambda < 0$. If $-1/2 \leq \sigma < 0$ then the roots are real-valued and have magnitude less than 1. If $\sigma < -1/2$ then the roots are complex-valued with magnitude equal to $(3 - 2\sigma)^{-1/2}$, which is less than 1. For this reason, it is impossible for the thresholding to be caused by an A-stability time-stepping constraint.

As an exploration of the thresholding of the time-steps, we performed simulations in which we use BDF2 with fixed small time steps, dt_{small}, until the solution has relaxed to steady state (up to round-off error) and then changed to larger, fixed time steps, dt_{large}, and continued to compute. The larger time steps were chosen to be larger than the threshold timestep, dt_∞, found by the adaptive time-stepper. The right plot of Figure 5 presents the local truncation error as a function of time. It appears to be growing exponentially fast in time. We have not yet discovered the cause of this instability. If one looks at graphs of the solutions, it appears that after the switch-over time, the value at $x = 0$, u_1 starts to increase from its near-steady-state value. This increase causes an increase at all values of v_i which then causes an increase at all values of u_i at the next time step. The value u_1 appears to be increasing faster than the

other values of u_i. When u_1 becomes larger than approximately 710, MATLAB evaluates $\exp(u_1)$ as "Inf" resulting in the BV-like flux constraint function $F(t)$ in (23) being evaluated as Inf. Solving for u and v at the next time step yields "NaN". If we define the time at which NaN is first observed as the blow-up time, we find that as the larger time-step, dt_{large}, decreases the "blow-up" time increases. This suggests that the thresholding in the time-step size is caused by the adaptive time-stepper's avoiding some sort of numerical instability at the $x = 0$ boundary. A related phenomenon is that when running the simulation with a fixed time step, we find that larger time steps result in slower relaxation[2] to steady-state. If the fixed time step is above the threshold time-step size, t_∞, the local truncation errors don't decrease — they increase.

3.2 The Effect of Boundary Conditions

It is natural to ask whether the thresholding behaviour still occurs when simpler (homogeneous, uncoupled) boundary conditions are used. First, we tested the VSSBDF2 solver on the toy model with homogeneous Dirichlet boundary conditions on the left boundary, $u(0,t) = 0$, and inhomogeneous, time-dependent Dirichlet boundary conditions at the right boundary, $u(1,t) = \exp(-t)$. We found that no thresholding occurs (see Figure 3.10 in [48]). This would seem to indicate that the thresholding does not occur simply due to solving a coupled parabolic-elliptic system of PDEs. Next, we tested the solver with various flux constraint functions in the boundary condition (24). We considered a) $u(0,t) = 0$ and $G(t) = 0$, b) $u(0,t) = 0$ and $G(t) = v(1,t) - 1/10$, and c) $u(0,t) = 0$ and $G(t) = u(1,t)\exp(v(1,t)) - \exp(-v(1,t))$. For all three sets of boundary conditions, we found that the time-step size thresholds (see Figure 3.11 in [48]).

These tests show that the thresholding comes about due to having the flux for u coupled to v and due to time stepping the boundary node. Even for the simplest BCs of $u(0,t) = 1$ and $G(t) = 0$ the term uv_x in the flux must be extrapolated, and so any flux boundary condition involves extrapolating a coupled term regardless of the flux constraint functions $F(t)$ and $G(t)$. It is strange to blame the thresholding on this extrapolation, however, since the uv_x term also gets extrapolated when time-stepping the values at interior points, but this extrapolation at interior nodes does not cause thresholding as demonstrated by the simulations with $u(0,t) = 0$ and $u(1,t) = \exp(-t)$. But the extrapolation seems to cause an issue at the boundary, even though we are time stepping the two boundary nodes exactly the same way as interior nodes.

Ultimately, both the "direct" method (27)–(28) and the "ghost point" approach (29)–(30) have time-steps that threshold but the ghost point approach is superior to the direct method in other ways and so it is clearly the better alternative.

[2]We say that the solution has "reached" steady state when the local truncation error is approximately 10^{-14}.

4 Numerical Tests on PNP-FBV Equations

Having studied the adaptive time stepper for the toy model (21)–(26), we now consider the problem of interest: the PNP-FBV equations. From the study of the toy model, we are confident that if the potential has voltage boundary conditions (8)–(9) then the time-stepper will be second-order accurate. For this reason, we consider current, rather than voltage boundary conditions. This requires time-stepping the additional ODE (10). Table 3 shows a convergence test using constant time steps. The scheme is at least second-order accurate, as expected.

dt	Ratios
0.005	4.43
0.005/2	4.23
0.005/4	4.44
0.005/8	4.71
0.005/16	4.94

Table 3: Convergence test on PNP-FBV equations (1)–(9) with current boundary conditions (10). The imposed external current is $j_{\text{ext}} = 0$. All model parameters were set to 1 except $\epsilon = 0.1$. See the caption of Table 1 for a full explanation. The mesh width is $dx = 1/300$ and the simulation ran to time $t = 0.1$. Runs with $dt = 0.001/32$ and $dt = 0.001/64$ were necessary to create the ratios.

We then considered the PNP-FBV equations with voltage boundary conditions. The voltage was chosen to be nearly piecewise constant in time; there are fast voltage increases at four times as shown in the left plot of Figure 6. Before the first voltage increase, the solution is relaxing to a steady state. The solution starts to relax to another steady state after each subsequent voltage increase. The time-steps chosen by the adaptive stepper are shown in the right plot of Figure 6. After a short transient, the adaptive time-stepper takes time steps that are near the threshold value and responds to each voltage step by sharply decreasing the time-step size followed by a coarsening of the step size. This is a demonstration of the value of adaptive time stepping versus constant time steps: a constant step method would require step sizes of 10^{-5} to resolve the imposed voltage $v(t)$ used here, whereas the adaptive stepper has step sizes larger than 10^{-2} for most of the duration of the simulation. The adaptiveness of the time-stepper allowed us to numerically explore certain standard experimental protocols, such as linear sweep voltammetry, in which a time-dependent voltage or current is applied [50].

Finally, we consider the singular perturbation parameter, ϵ, which controls the dynamics of the boundary layers near the electrodes in the PNP-FBV equations (2). Figure 7 presents some steady state solutions for PNP-FBV equations (1)–(9) with current boundary conditions (10). The figure on the left presents the steady-state concentration of the positive species, $c_+(x)$, for three values of ϵ. The smaller the value of ϵ, the thinner the transition layer near $x = 0$ and $x = 1$ is. The figure on the right presents the corresponding steady-state potential, $\phi(x)$.

For simulations in which the solutions are converging to a steady-state solution, we find that the time step size thresholds to a limiting value, dt_∞. Interestingly, compared to the behaviour for adaptive

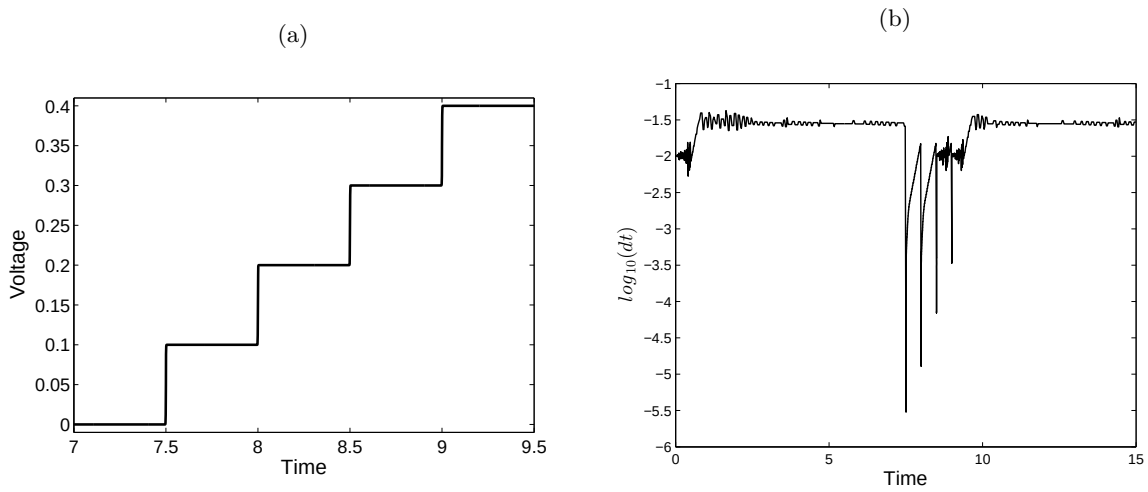

Figure 6: PNP-FBV equations (1)–(9) with $N = 90$, tol = 1e-6, $range = tol/3$, all model parameters set to 1 except $\epsilon = 0.1$. The initial data for the concentrations is $c_{\pm}(x, 0) = 1 + .1\sin(2\pi x)$. The left plot shows the time-dependent voltage $v(t)$. The voltage values are nearly 0, 0.1, 0.2, 0.3 and 0.4 with sharp transitions at times 7.5, 8, 8.5 and 9. The voltage is constructed using a sum of $1/20\,(\tanh(1000(t-t_0))+1)$ functions with appropriately chosen values of t_0. The right plot presents the time-step sizes as a function of time.

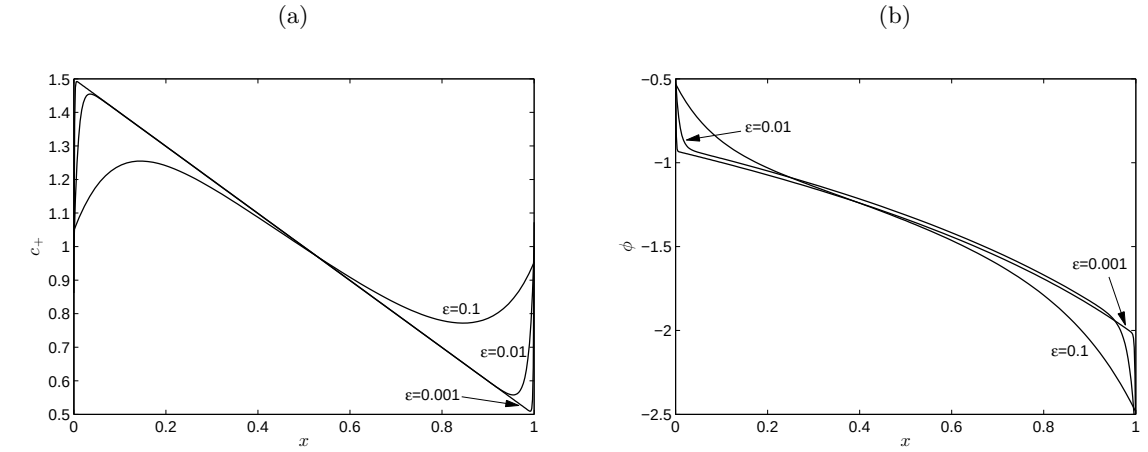

Figure 7: Steady-state solutions for PNP-FBV equations (1)–(9) with current boundary conditions (10). The imposed external current is $j_{\text{ext}} = 0.5$. The simulations were run until time $t = 5$; the profiles at that time are shown here. Three values of ϵ were used. For $\epsilon = 0.1, 0.01$, a uniform-in-space mesh was used with $dx = 1/90$. For $\epsilon = 0.001$, a piecewise uniform-in-space mesh was used with $dx = 1/600$ in $[0, 0.1]$ and $[0.9, 1]$ and $dx = 1/75$ in $[0.1, 0.9]$. Left figure: concentration of positive species, c_+, versus x. Right figure: potential, ϕ, versus x.

time-stepping the diffusion equation with Forward Euler (see Figure 4) for which dt_∞ depends on the meshwidth, we find for the PNP-FBV equations that dt_∞ seems to depend only on ϵ. We find that if one decreases ϵ by a factor of a then the limiting time-step size decreases by a^2; see Table 4. The PNP equations (1)–(2) have been nondimensionalized. In doing this, a nondimensional time scale is introduced; it scales like $D\tau/L^2$ where D is the diffusivity, τ is the dimensional time unit, and L is the width of the cell. Also, $\epsilon = \lambda_D/L$ where λ_D is the Debye screening length. As a result, ϵ scales with $1/L$ and t (and therefore dt) scales with $1/L^2$. This explains why lowering epsilon by a factor of a should result in dt being smaller by a factor of a^2.

ϵ	dt_∞	dt_∞/ϵ^2
0.1	0.0257	2.57
0.01	0.000241	2.41
0.001	3.70×10^{-6}	3.70
10^{-4}	2.48×10^{-8}	2.48
10^{-5}	2.25×10^{-10}	2.25
10^{-6}	2.37×10^{-12}	2.37

Table 4: The PNP-FBV equations (1)–(9) are simulated with the applied voltage $v(t) = 0$. Six values of ϵ are considered. We denote the limiting timestep size chosen by the adaptive stepper by dt_∞. We compute dt_∞ by averaging the time-step sizes, dt, over a time interval of length 1 chosen late in the simulation.

In this way, by taking small enough time steps, our adaptive time-stepper is able to run with small values of ϵ. This is an important contribution, as previous work on time-dependent solutions to the PNP-FBV equations [22, 51] took comparably large values, $\epsilon \geq 10^{-3}$, whereas in principle we are able to simulate with any value of ϵ.

5 Conclusions and Future Work

In this work, we considered the Poisson-Nernst-Planck equations with generalized Frumkin-Butler-Volmer reaction kinetics at the electrodes. We developed a solver that dynamically chooses the time-step size so that the local truncation error is within user-specified bounds. The spatial discretization can be nonuniform, allowing for the finer mesh needed near the electrodes to resolve the boundary layers. The adaptive time-stepping allows the solver to adaptively choose small time steps during the initial transient period during which the boundary layers may be forming quickly in response to the boundary conditions. The adaptive time-stepping also allows one to study physical situations in which the imposed voltage or current have sudden, fast changes. We paid careful attention to the rate of convergence in time and to the handling of boundary conditions. The numerical method was tested on a simplified "toy" model which retains some key features of the full model, and evidence for a nonlinear instability was found which appears to originate from the boundary conditions. Furthermore, for the PNP-FBV equations our numerical scheme was found to be able to operate with (arbitrarily) small values of the nondimensional parameter ϵ, although the smaller the value of ϵ the smaller the resulting allowable time steps were and hence the longer the simulation time was.

For future work, the nonlinear instability at the boundary is worth investigating further, so as to determine the exact cause and potentially modify the scheme to remove it. Also, our method only automatically adjusts time steps; the spatial mesh is chosen before the simulation begins. Incorporating adaptive mesh refinement along with adaptive time stepping would further improve on the numerical method.

A Derivation of the Local Truncation Error Formula and the Extrapolation Formula

In this appendix, the error approximation and extrapolation formula used in Subsection 2.3 is derived. We consider time stepping the ODE

$$\frac{du}{dt} = f(u) + g(u) \tag{37}$$

where the value of u at time $n+1$ is given by

$$\frac{1}{dt_{\text{now}}}\left(\frac{1+2\omega}{1+\omega}u^{n+1} - (1+\omega)u^n + \frac{\omega^2}{1+\omega}u^{n-1}\right) = (1+\omega)f(u^n) - \omega f(u^{n-1}) + g(u^{n+1}) \tag{38}$$

By making the substitutions $u^{n+1} = u(t+dt_{\text{now}})$ and $u^{n-1} = u(t-dt_{\text{old}})$ and performing a multivariable Taylor expansion about t, the local truncation error is found to take the form

$$\begin{aligned} LTE = {} & \left(u'(t) - f(u(t)) - g(u(t))\right) dt_{\text{now}} + \left(u''(t) - f'(u(t))\, u'(t) - g'(u(t))\, u'(t)\right) dt_{\text{now}}^2 \\ & + \left(\frac{1}{2}f''(u(t))\, u'(t)^2 + \frac{1}{2}f'(u(t))\, u''(t) - \frac{1}{6}u^{(3)}(t)\right) dt_{\text{old}}\, dt_{\text{now}}^2 \\ & + \left(-\frac{1}{2}g''(u(t))\, u'(t)^2 - \frac{1}{2}g'(u(t))\, u''(t) + \frac{1}{3}u^{(3)}(t)\right) dt_{\text{now}}^3 + O(dt^4) \end{aligned} \tag{39}$$

where $O(dt^4)$ denotes terms where dt_{now} and dt_{old} combined appear four or more times. If $u(t)$ is a thrice-differentiable solution of (37), the dt_{now} and dt_{now}^2 terms vanish and the cubic term can be simplified, resulting in

$$\begin{aligned} LTE = {} & \left(\frac{1}{2}f''(u(t))\, u'(t)^2 + \frac{1}{2}f'(u(t))\, u''(t) - \frac{1}{6}u^{(3)}(t)\right) dt_{\text{old}} dt_{\text{now}}^2 \\ & + \left(-\frac{1}{2}g''(u(t))\, u'(t)^2 - \frac{1}{2}g'(u(t))u''(t) + \frac{1}{3}u^{(3)}(t)\right) dt_{\text{now}}^3 + O(dt^4) \end{aligned} \tag{40}$$

$$= \left(\frac{1}{3}u^{(3)}(t) - \frac{1}{2}g''(u(t))\, u'(t)^2 - \frac{1}{2}g'(u(t))\, u''(t)\right) dt_{\text{now}}^2\, (dt_{\text{now}} + dt_{\text{old}}) + O(dt^4). \tag{41}$$

This means that if $u^{(3)}(t)$ is bounded, the local truncation error for the coarse step is

$$\epsilon_c = u_c^{n+1} - u(t^{n+1}) \approx C dt_{now}^2\, (dt_{now} + dt_{old}) \tag{42}$$

To find the local truncation error for u_f^{n+1}, we find the local truncation error for $u_f^{n+\frac{1}{2}}$ by replacing dt_{now} and dt_{old} in (42) by $\frac{dt_{\text{now}}}{2}$ and $\frac{dt_{\text{old}}}{2}$, respectively. We then find the local truncation error for u_f^{n+1} by setting both dt_{now} and dt_{old} to $\frac{dt_{\text{now}}}{2}$. Adding these local truncation errors yields the LTE for u_f^{n+1},

$$\epsilon_f = u_f^{n+1} - u(t^{n+1}) \approx \frac{C dt_{now}^2\, (dt_{now} + dt_{old})}{8} + C\, \frac{dt_{now}^3}{4} \tag{43}$$

Using equations (42) and (43), we can approximate C using u_c, u_f, dt_{now} and dt_{old} and then use this approximation in (42), resulting in

$$\epsilon_c \approx \frac{u_c - u_f}{\frac{5}{8}dt_{now}^3 + \frac{7}{8}dt_{now}^2 dt_{old}}\, dt_{now}^2\,(dt_{now} + dt_{old}) \tag{44}$$

Finally, we subtract ϵ_c from u_c^{n+1} to create a more accurate approximation of $u(t^{n+1})$. This results in the extrapolation formula $u(t^{n+1}) \approx u_c^{n+1} - \epsilon_c = \alpha u_c^{n+1} + \beta u_f^{n+1}$ where

$$\alpha = -\frac{dt_{old} + 3\,dt_{now}}{7\,dt_{old} + 5\,dt_{now}}, \quad \beta = 8\,\frac{dt_{old} + dt_{now}}{7\,dt_{old} + 5\,dt_{now}}. \tag{45}$$

B Acknowledgements

We thank Steven Ruuth for helpful conversations and encouragement.

References

[1] P. M. Biesheuvel and M. Z. Bazant. Nonlinear dynamics of capacitive charging and desalination by porous electrodes. *Phys. Rev. E*, 81(3):031502, 2010.

[2] P. M. Biesheuvel, Yeqing Fu, and M. Z. Bazant. Diffuse charge and Faradaic reactions in porous electrodes. *Phys. Rev. E*, 83(6):061507, 2011.

[3] P. M. Biesheuvel, Yeqing Fu, and M. Z. Bazant. Electrochemistry and capacitive charging of porous electrodes in asymmetric multicomponent electrolytes. *Russian J. Electrochem.*, 48(6):580–591, 2012.

[4] P. B. Peters, R. van Roij, Martin Z. Bazant, and P. M. Biesheuvel. Analysis of electrolyte transport through charged nanopores. *Phys. Rev. E*, 93:053108, 2016.

[5] Ian Streeter and Richard G. Compton. Numerical simulation of potential step chronoamperometry at low concentrations of supporting electrolyte. *J. Phys. Chem. C*, 112(35):13716–13728, 2008.

[6] Richard G. Compton and Craig E. Banks. *Understanding Voltammetry.* Imperial College Press, London, 2011.

[7] E. Victoria Dydek and Martin Z. Bazant. Nonlinear dynamics of ion concentration polarization in porous media: The leaky membrane model. *AlChE Journal*, 59(9):3539–3555, 2013.

[8] Victor V. Nikonenko, Natalia D. Pismenskaya, Elena I. Belova, Philippe Sistat, Patrice Huguet, Gérald Pourcelly, and Christian Larchet. Intensive current transfer in membrane systems: Modelling, mechanisms and application in electrodialysis. *Adv. Colloid Interface Sci.*, 160(1):101–123, 2010.

[9] A. Yaroshchuk. Over-limiting currents and deionization shocks in current-induced polarization: local equilibrium analysis. *Adv. Colloid Interface Sci.*, 183:68–81, 2012.

[10] Martin Z. Bazant and Todd M. Squires. Induced-charge electrokinetic phenomena. *Current Opinion in Colloid & Interface Science*, 15(3):203–213, 2010.

[11] Martin Z. Bazant, Mustafa Sabri Kilic, Brian D. Storey, and Armand Ajdari. Towards an understanding of induced-charge electrokinetics at large applied voltages in concentrated solutions. *Advances in colloid and interfacial science*, 152(1):48–88, 2009.

[12] Martin Z. Bazant, Brian D. Storey, and Alexei A. Kornyshev. Double layer in ionic liquids: Overscreening versus crowding. *Phys. Rev. Lett.*, 106(4):046102, 2011.

[13] Alexei A. Kornyshev. Double-layer in ionic liquids: Paradigm change? *J. Phys. Chem. B*, 111:5545–5557, 2007.

[14] Martin Z. Bazant, Kevin T. Chu, and B. J. Bayly. Current-voltage relations for electrochemical thin films. *SIAM J. Appl. Math.*, 65(5):1463–1484, 2005.

[15] Kevin T. Chu and Martin Z. Bazant. Electrochemical thin films at and above the classical limiting current. *SIAM J. Appl. Math.*, 65(5):1485–1505, 2005.

[16] P. M. Biesheuvel, M. van Soestbergen, and M. Z. Bazant. Imposed currents in galvanic cells. *Electrochemica Acta*, 54:4857–4871, 2009.

[17] P. Maarten Biesheuvel, Alejandro A. Franco, and Martin Z. Bazant. Diffuse charge effects in fuel cell membranes. *J. Electrochem. Soc.*, 156(2):B225–B233, 2009.

[18] Alpha A Lee, Svyatoslav Kondrat, Gleb Oshanin, and Alexei A Kornyshev. Charging dynamics of supercapacitors with narrow cylindrical nanopores. *Nanotechnology*, 25(31):315401, 2014.

[19] Martin Z. Bazant. Theory of chemical kinetics and charge transfer based on nonequilibrium thermodynamics. *Accounts of Chemical Research*, 46(5):1144–1160, 2013.

[20] M. van Soestbergen. Frumkin-Butler-Volmer theory and mass transfer in electrochemical cells. *Russian Journal of Electrochemistry*, 48(6):570–579, 2012.

[21] A. A. Moya, J. Castilla, and J. Horno. Ionic transport in electrochemical cells including electrical double-layer effects. A network thermodynamics approach. *J. Phys. Chem.*, 99:1292–1298, 1995.

[22] M. van Soestbergen, P. M. Biesheuvel, and M. Z. Bazant. Diffuse-charge effects on the transient response of electrochemical cells. *Phys. Rev. E*, 81(2):1–13, 2010.

[23] John Newman and William H. Smyrl. The polarized diffuse double layer. *Trans. Faraday Soc.*, 61:2229–2237, 1965.

[24] William H. Smyrl and John Newman. Double layer structure at the limiting current. *Trans. Faraday Soc.*, 63:207–216, 1966.

[25] Martin Z. Bazant, Katsuyo Thornton, and Armand Ajdari. Diffuse-charge dynamics in electrochemical systems. *Phys. Rev. E*, 70(2):021506, 2004.

[26] H. Cohen and J. W. Cooley. The numerical solution of the time-dependent Nernst-Planck equations. *Biophys. J.*, 5(2):145–162, 1965.

[27] James R. Sandifer and Richard P. Buck. An algorithm for simulation of transient and alternating current electrical properties of conducting membranes, junctions and one-dimensional, finite galvanic cells. *J. Phys. Chem.*, 79(4):384–392, 1975.

[28] Timothy R. Brumleve and Richard P. Buck. Numerical solution of the Nernst-Planck and Poisson equation system with applications to membrane electrochemistry and solid state physics. *J. Electroanal. Chem. Interfacial Electrochem.*, 90(1):1–31, 1978.

[29] Tomasz Sokalski, Peter Lingenfelter, and Andrzej Lewenstam. Numerical solution of the coupled Nernst-Planck and Poisson equations for liquid junction and ion selective membrane potentials. *J. Phys. Chem. B*, 107(11):2443–2452, 2003.

[30] W. D. Murphy, J. A. Manzanares, S. Mafé, and H. Reiss. Numerical simulation of the nonequilibrium diffuse double layer in ion-exchange membranes. *J. Phys. Chem.*, 97(32), 1993.

[31] G. W. Gear. *Numerical Initial Value Problems in Ordinary Differential Equations.* Prentice Hall, Englewood Cliffs, NJ, 1971.

[32] Donald L. Scharfetter and Herman K. Gummel. Large-signal analysis of a silicon read diode oscillator. *IEEE Trans. Electron Devices*, 16(1):64–77, 1969.

[33] Silvano Erlicher, Luca Bonaventura, and Oreste S. Bursi. The analysis of the generalized- method for non-linear dynamic problems. *Comput. Mechanics*, 28(2):83–104, 2002.

[34] Dieter Britz and Jörg Strutwolf. Several ways to simulate time dependent liquid junction potentials by finite differences. *Electrochimica Acta*, 137:328–335, 2014.

[35] Dieter Britz. *Digital Simulation in Electrochemistry.* Springer, Berlin, 2005.

[36] Richard G. Compton, Eduardo Laborda, and Kristopher R. Ward. *Understanding Voltammetry: Simulation of Electrode Processes.* Imperial College Press, London, 2014.

[37] D. Wang and S. J. Ruuth. Variable step-size implicit-explicit linear multistep methods for time-dependent partial differential equations. *Journal of Computational Mathematics*, 26(6):838–855, 2008.

[38] O. Axelsson, X. He, and M. Neytcheva. Numerical solution of the time-dependent Navier-Stokes equation for variable density-variable viscosity. part I. Technical report, Uppsala University, 2014.

[39] Daniel Lecoanet, Josiah Schwab, Eliot Quataert, Lars Bildsten, F. X. Timmes, Keaton J. Burns, Geoffrey M. Vasil, Jeffrey S. Oishi, and Benjamin P. Brown. Turbulent chemical diffusion in convectively bounded carbon flames. *Astrophys. J.*, 832(1):71, 2016.

[40] Gunilla Linde, Jonas Persson, and Lina Von Sydow. A highly accurate adaptive finite difference solver for the Black-Scholes equation. *Int. J. Comp. Math.*, 86(12):2104–2121, 2009.

[41] S. Eckert, H. Baaser, D. Gross, and O. Scherf. A BDF2 integration method with step size control for elasto-plasticity. *Comput. Mech.*, 34(5):377–386, 2004.

[42] J. Rosam, P. K. Jimack, and A. Mullis. A fully implicit, fully adaptive time and space discretisation method for phase-field simulation of binary alloy solidification. *J. Comput. Phys.*, 225(2):1271–1287, 2007.

[43] Lloyd N. Trefethen. *Approximation Theory and Approximation Practice*. Siam, 2013.

[44] Rodrigo B. Platte and Lloyd N. Trefethen. Chebfun: a new kind of numerical computing. In *Progress in Industrial Mathematics at ECMI 2008*, pages 69–87, 2008.

[45] Roger Peyret. *Spectral Methods for Incompressible Viscous Flow*. Springer-Verlagl, New York, 2002.

[46] U. M. Ascher, S. J. Ruuth, and B. T. R. Wetton. Implicit-explicit methods for time-dependent partial differential equations. *SIAM Journal on Numerical Analysis*, 32(3):797–823, 1995.

[47] Ernst Hairer, Syvert P Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems (Springer Series In Computational Mathematics)*. Springer Berlin, Heidelberg, 2009.

[48] David Yan. *Macroscopic Modeling of a One-Dimensional Electrochemical Cell Using the Poisson-Nernst-Planck Equations*. PhD thesis, University of Toronto, 2017.

[49] J. M. Thomas. *Numerical Partial Differential Equations: Finite Difference Methods*. Springer, New York, 1995.

[50] David Yan, Martin Z. Bazant, P. M. Biesheuvel, Mary C. Pugh, and Francis P. Dawson. Theory of linear sweep voltammetry with diffuse charge: Unsupported electrolytes, thin films, and leaky membranes. *Phys. Rev. E*, 95(3):033303, 2017.

[51] Laurits Højgaard Olesen, Martin Z. Bazant, and Henrik Bruus. Strongly nonlinear dynamics of electrolytes in large ac voltages. *Phys. Rev. E*, 82(1):011501, 2010.